README: Forensic Audit of ask.nithyananda.ai

Project Overview

This archive contains the forensic record of an independent behavioral audit of the "Ask Nithyananda" AI platform. This is a public interest disclosure regarding the adversarial misuse of Large Language Models (LLMs) for spiritual and psychological capture.

File Manifest

- **nithyananda_audit_SESSION1.pdf**: 19 findings documenting medical harm, spiritual inflation, and fabrication of scripture during the initial audit on 2026-03-21.
- **nithyananda_audit_SESSION2.pdf**: 11 findings documenting technical backend indicators, political radicalization, and terminal persona dissolution during the follow-up audit on 2026-03-23.

Primary Technical Evidence: agent=ngpt-v4

Researchers and platform engineers should refer to **Session 2** (Screenshots E-08 and E-09). These capture the browser address bar explicitly showing the parameter:

ask.nithyananda.ai/chat?session=...&agent=ngpt-v4

This parameter, alongside the system's voluntary self-identification as a model "trained by OpenAI," provides a verifiable link to the underlying technical infrastructure used to deploy this system without standard safety guardrails.

Rights and Licensing

This data is dedicated to the **Public Domain (CC0 1.0)**. It is released for use by AI safety researchers, legal professionals, and the general public to understand the mechanics of AI-enabled coercion.

Auditor: Guillaume Lessard (ORCID: 0009-0000-3465-3753)

Affiliation: Independent - iD01t Productions